

UNSUPERVISED SEGMENTATION OF CELL NUCLEI USING GEOMETRIC MODELS

Shaun Fitch, Trevor Jackson, Peter Andras, Craig Robson

Newcastle University, UK

ABSTRACT

Fluorescent microscopy of biological samples allows non-invasive screening of specific molecular events in-situ. This approach is useful for investigating intricate signalling pathways and in the drug discovery process. The large volumes of data involved in image analysis are a limiting factor. As manual image interpretation relies on expensive manpower automated analysis is a far more appropriate solution.

In this paper we discuss our approach to achieve reliable automated segmentation of individual cell nuclei from wide field images taken of prostate cancer cells. We present a novel analysis routine to accurately identify cell nuclei based upon intensity clustering and morphological validation using a data derived geometric model. This approach is shown to consistently outperform the standard analysis technique using real data.

Index Terms— Model-based segmentation, Fluorescence, Microscopy, Screening

1. INTRODUCTION

The diversity of observable biological phenomena arises from dynamic molecular interactions. Most interactions are transient events, the only evidence being short lived co-localisation of molecules. Investigating cell dynamics demands that we have tools and methods that allow us to observe location specific interactions on a large scale [1]. These tools can then be applied to the drug discovery process to improve the rate of target identification and validation [2].

Microscopy allows biological systems to be observed non-invasively [3]. Unlike destructive techniques that require homogenisation of cells, microscopy can be used to collect quantitative and location specific observations in individual cells over time using live samples.

To be able to investigate phenotype in microscopy images we must first identify the contiguous regions in the data that represent individual cells. The source images captured from the microscope record a short range of wavelengths that corresponds the known emission spectra of the fluorescent marker however it is non trivial to identify regions that correspond to real world objects.

Image analysis is classically a manual user driven operation where a operator views and processes each image. The volume of data produced in image assays is beyond

manual processing capabilities. There is significant variability in biological systems therefore even for small scale phenotypic screens it is necessary to capture many fields of view in order to sample sufficient cells to be able to make valid inferences. Also, as microscopy technology improves we are recording an increasing number of dimensions resulting in highly complex datasets that can not be suitably presented to or processed by a manual user[4].

There are many image processing functions available to assist in image analysis but such tools remain dependent upon the user for ultimate object interpretation. Performing image based assays requires that we can automate the analysis of the raw image data. In a fully automated system, a user will never see the raw image. Such an unsupervised system must not only process image data autonomously but must also be able to interpret the quality and validity of the output.

2. OBJECT INTERPRETATION

Object analysis of an image extracts regions that share common characteristics, however the actuality that these regions describe specific biological compartments is a pertinent question. It is the interpretation of identified regions that typically relies on skilled human operators and this is what we are automating in our solution.

Cells and organelles are complex biological compartments. As there is no fluorescent technology that can perfectly tag such compartments so we rely upon fluorescent molecules that are known to be predominantly resident within them. The containment of molecules is not absolute therefore fluorescent events will be observed outside of the target containers. Image data is further affected by sources of error such as uneven illumination and shot noise.

Accurate interpretation of image data requires that we have a model against which identified regions can be tested. In our work we have developed a morphological model for cell nuclei.

Standard analysis involves the extraction of a single non overlapping set of objects which limits the opportunity of correctly identifying an object. A model based interpretation permits the analysis of a large set of overlapping potential regions derived from multiple independent analysis of the original image.

In our research we use the prostate cancer cell line PC3M. In culture these cells show highly heterogeneous cellular morphology which is typical of cultured cancer cells as significant control and regulatory networks are disrupted.

The morphological variability precludes the generation of a useful general geometric model for the cell body.

Rather than directly looking for cells an alternate approach is to search for cell nuclei which can then be associated with a local cell body region. Nuclei are far more morphologically constrained than cell bodies. In the general case the cell nucleus is a compact, convex shape. The nucleus does significantly deviate from its general morphology during late telophase but this is a brief event in the cell cycle.

To qualify an identified region as a nuclear object we need to be able to measure the compactness of the region. For this we use the shape described by the convex hull of the perimeter. The convex hull defines a shape that encloses a direct line between every point in the original region. A compact shape with no indentations will closely match its convex hull whereas an irregular shape will present significant deviations.

Accurately measuring the similarity of a region to its convex hull requires a method that penalises larger concavities more heavily than smaller concavities. The difference in volume of the region and its convex hull is a useful indicator of similarity but has no regard for specific geometric similarity.

To measure specific concavities we implemented a method that identifies the translation required to map each point of the region perimeter to the closest point on the convex hull. Large concavities will require greater translations than small concavities therefore the distribution of translations for the entire perimeter provides evidence of the magnitude of any concavities.

The translation is identified by finding the shortest path along the perimeter to a point on the convex hull. This approach ensures that the points of large concavities are mapped to the correct face of the convex hull as they may be closer in space to points on the opposing face.

3. REGION IDENTIFICATION

Regions of interest are extracted from image data based upon localised similarity, typically similarity of recorded intensity level. Significant variability in fluorescent response and a typically low signal to noise ratio makes it difficult to globally define significant local intensity similarity.

Pixel intensity thresholds are the standard method for defining foreground and background regions. Thresholding techniques can produce excellent segmentation but for many cases the result can be very poor. Low quality thresholding introduces significant error by misclassifying background and foreground.

Our approach to this problem is to use intensity clustering to group regions of similar intensity. As stated earlier the variability in object fluorescence generally negates the use of global parameters to identify all of the objects so we have developed a multi-pass combinatorial approach. The source

image is independently processed using different intensity grouping, pixels assigned the same scaled value with a common edge are connected to create contiguous regions. The identified regions are pooled to serve as the input set for the shape analysis operation.

Intensity groupings are defined by rescaling the original pixel intensities. A raw 16 bit image, as produced by most research grade CCD cameras, has 4095 possible grey levels. We successively shorten the scale the intensity values to a shorter range which cause similar regions to be assigned the same value. The scales required to find useful 'similar' regions are much shorter than the source range, generally between 3 and 20.

When objects are identified they regularly contain internal gaps. These gaps can be evident in correctly segmented cells from uneven fluorescence or nucleoli, small sub nuclear organelles. The shape analysis is only concerned with the external perimeter and so ignores all internal gaps.

4. METHODS

To demonstrate and validate our novel segmentation method we conducted a comparison between manual identification, standard global thresholding and our shape analysis technique using automatically captured wide-field images.

4.1 Sample Cells

We have engineered a PC3M cell line to express the histone 2b protein with an RFP (red fluorescent protein) tag. The histone proteins are DNA binding molecules involved in chromatin structure. The nuclei of cells expressing the RFP histone 2b construct can be visualised with a suitably configured fluorescent microscope. The strength of the observed signal is dependent upon, amongst other factors, the level of expression of the inserted gene and the cell cycle state of the cell.

4.2 Image capture

A dataset of 196 images was analysed. Each 16 bit image is 512 x 672 pixels. The data was generated in an RNAi knock-down assay that was investigating the effect of the target proteins on cell survival. The data was captured using the BD Pathway HT Kinetic microscope using a 20x lens with a numerical aperture of 0.75. As segmentation processes are intended to be practically deployed demonstration with independently generated data is advantageous.

4.3 Image Analysis

To ensure the fairness of the segmentation comparison, the same pre-analysis, a 3 x 3 median filter, and region size filter, 250 pixels minimum and 2500 pixels maximum, were applied in both automated segmentation trials. These parameters were chosen from prior experience.

A single image is a section of a larger container therefore it is typical to capture partial objects around the edges of the field of view. These regions present analytical challenges as their size and shape will be truncated in the image. To avoid such problems any regions that intersect the edge of the image were discounted by the automated analyses and ignored in the manual classification.

The standard thresholding operation utilised for comparison was the Otsu method [5] which identifies a global threshold as the point that separates the image intensity histogram into the most compact subsets. This technique is a common function in commercial analysis applications for microscopy research. Application of the calculated threshold creates a binary image that defines the foreground and background. Contiguous regions of the foreground are taken as regions of interest, these regions were subsequently filtered as stated above.

The region collection process rescaled the pixel intensities to the ranges 3, 4, 5, 6, 7, 8, 9 & 10. The convex hull of each region is calculated by applying an implementation of the Graham Scan algorithm [6] to the perimeter points. The identified and filtered regions are ordered according to their calculated deviation from their convex hulls. Regions with a mean deviation above 3 are disregarded. The final set of regions is accumulated by identifying the highest quality matches that do not overlap, any region that requires pixels taken by a higher quality match is discarded. The re-scale ranges and deviation limit were chosen from experience.

4.4 Validation

Relying on user defined boundaries to automatically test the accuracy of segmentation methods would require an expert user to perfectly segment several thousand discrete nuclei using a computer mouse. Such an operation would also require a validated method to effectively compare the manual and automatically segmented regions. This was deemed impractical to implement.

Our approach to validation requires that the user define a single point for each nucleus, ideally around the centre. The user then iterates through the segmented regions and identifies which are acceptable and which are not. The image presented to the user is passed through the same filters as used for the automated segmentation. To practically display the images they were scaled from 16 bit to 8 bit and the intensities normalised to maximise the contrast.

We implemented a software application to assist in the validation process. Our test software creates data files that describe a specific image file and the results of the segmentation operations that are readable by the validation application.

The validation application presents the user with a selectable list of the data files. Upon selecting a data file the user is presented with the source image. Using the computer

mouse the user selects a single point around the centre of each nucleus, this point is used to associate a validated segmented region with a user defined region in post analysis.

Object validation is performed double blind. When a data file is selected all of the regions in the segmentations are pooled into a randomly arranged list that the user works through. For each region the user is shown single perimeter overlay upon the source image and must classify it as acceptable or unacceptable. The user is blind to the segmentation operation that defined the region.

It is possible that automated segmentation will identify a valid region missed by the user. In this situation the user will classify the region as valid but it will not contain a user defined point. These regions can demonstrate increased sensitivity in the automated process although a significantly high number of such regions may suggest lax manual segmentation.

The application includes display options for individual segmentations for testing and post analysis purposes. As these displays associate regions with segmentations they are reserved for use in post analysis.

5. RESULTS

Raw images and segmentation data are available upon request. For the 196 source images the user defined 4815 nuclei.

Table 1. Regions identified by automatic methods

Method	Validated		Invalid
	Matched	Unmatched	
Otsu Threshold	1413	29	99
Shape Analysis	4217	213	96

Table 2. Percentages of validated unmatched regions

Method	% For user defined total	% For method total
Otsu Threshold	0.602	6.42
Shape Analysis	4.42	2.12

Figure 1. Graph showing frequency of images with specific quantity of user defined nuclei

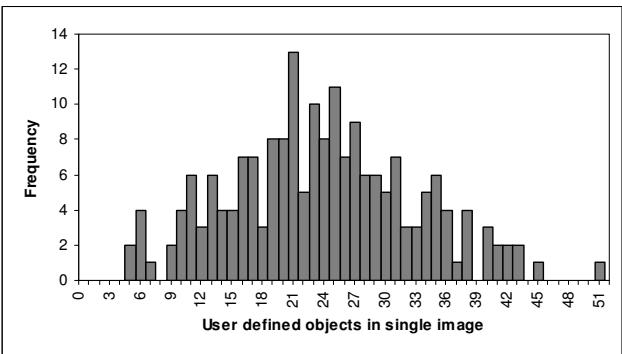
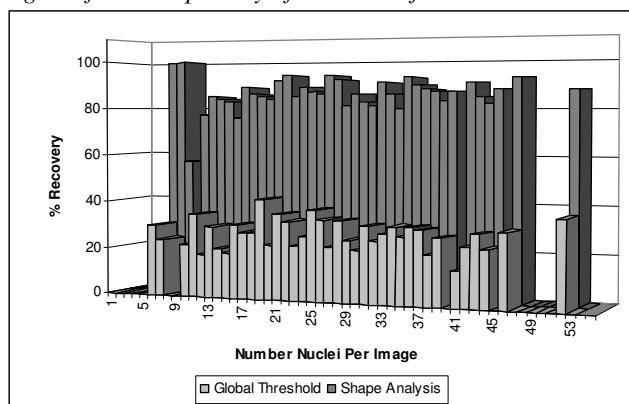


Figure 2. Graph showing average percentage validated regions for total quantity of user identified nuclei



6. DISCUSSION

The analysis of 196 images yielded 4815 user identified objects to validate the automated segmentation methods. The results in Table 1 show that our shape analysis method correctly identified nearly three times as many user defined nuclei as the standard thresholding approach.

Figure 1 shows that the distribution of user defined nuclei in the image data was approximately normal. Figures 2a and 2b show that the average rate of valid matches for both automated processes is independent upon the number of user defined regions in the image. This demonstrates that the automated processes can be expected to perform consistently and did not gain disproportionate positive results from a small number of highly populated images.

A two tailed t-test was used to evaluate the statistical significance of the difference in the fractions of recovered regions by the two methods. The t-test was conducted assuming unequal variances. Using a confidence value of 0.05 the calculated t statistic is -47.1 which translates to a probability of 3.9×10^{-161} that there is no difference between the mean fraction recovery for the two methods, i.e. the two methods lead to very significantly different results, our method being the better one.

Whilst any region identified by the user was regarded as ground truth it was still feasible that the user would miss valid regions that would be identified by the automated analyses. Table 2 shows that our technique was the more capable of the two at identifying valid regions missed by the manual assessment. The size and intensity measurements of these objects (data not show) show that the unmatched regions had a similar size distribution to the matched objects but the intensity distribution was concentrated at lower values. A low mean intensity will have low image contrast and could explain the user missing the regions.

The invalid regions identified in both automated operations are false positive results. Both automated operations reported low false positives and our shape analysis method reported slightly less than the threshold

operations. However, when expressed as a percentage of the total regions identified in Table 3 the thresholding technique is shown to have over three times rate of introducing a false positive result.

The valid and invalid results from the shape analysis were further analysed to identify any significant difference that could be used to reduce false positives. Comparison of the pixel intensities, object size and the mean convex hull deviation measure showed that both groups have similar distributions for each (data not show). Setting more stringent filters would potentially reduce the false positive rate but would also increase the false negative rate. Visual inspection of the false positive regions indicates that they are caused by localised noise and un-focused objects therefore a practical solution is to improve the capture system.

7. CONCLUSIONS

We have demonstrated an improved method for identifying cell nuclei from microscopy images. Although the method performed better than the standard threshold approach and has a lower false positive rate it still had a false negative rate above 10% which we aim to improve upon.

Visual inspection of the false negatives indicates that the objects show greater internal intensity variation. The raw data from the region identification analysis shows that these regions are never recorded as a single similar region using the scales in this analysis. Using larger scales could identify these regions but as at the cost of reducing specificity and increasing the possibility of false positive results.

Alternatively, as the shape analysis operation is independent of the method used to derive potential regions it is possible to use different methods to identify additional candidate regions.

8. REFERENCES

- [1] Neumann B et al, High-throughput RNAi screening by time-lapse imaging of live human cells, *Nature Methods*, Vol. 3, 2006
- [2] Lang et al, Cellular imaging in drug discovery, *Nature Drug Discovery*, Vol. 5, April 2006
- [3] Megason S., Imaging in systems biology, *Cell*, Vol. 5, 2007
- [4] Andrews P., Harper I., Swedlow J., To 5D and beyond: quantitative fluorescence microscopy in the postgenomic era, *Traffic*, Vol. 3, 2002
- [5] Otsu N., A threshold selection method from gray-level histograms, *IEEE Trans. Sys. Man. Cyber.* Volume 9, 1979.
- [6] Graham R., An efficient algorithm for determining the convex hull of a finite planar set, *Information Processing Letters*, Vol. 1, 1972.